

**Compact Disc Write Once (CD-WO) or
Compact Disc Recordable (CD-R)
Volume and File Structure Standard
(aka Frankfurt Group Proposal)**

1. Introduction

Compact Disc - Read Only Memory ((3)-ROM) has fulfilled its promise of becoming a dominant data storage and distribution media of the 90's. Starting in October, all Macintosh computers will be shipped with CD-readers, every Multi-media PC (M PC) includes a drive, and nearly every UNIX workstation comes with a CD-reader for system software installation.

A key aspect of the success of the compact disc industry is ISO-9660:88, the international standard for the CD-ROM logical format. This standard allows the same CD-ROM to be read and interpreted on Macintosh, MS-1) OS, Unix VA X/V MS, and many other computer platforms. However, several design problems surfaced as ISO-9660 drivers were implemented on various operating systems. For example, the information needed to do a UNIX directory listing command is stored not in the directory records, but in Extended Attribute Records which are located with the data file. 'I'lllls to do a simple directory command, seeks to the individual data file locations must be done.

A major disadvantage of CD-ROM has been the need to go through a moderately expensive mastering process at a vendor's plant even for small numbers of replication and the long turn around time.. The emergence in the early 1990's of recordable CD's has allowed in-house production of disks. While recorder costs were prohibitive as little as two years ago (\$30,000), there are now several vendors offering recorders in the \$4,000 to \$6000 range, making a CD-recorder affordable at a department level, if not on an individual level at this time. The drops in recorder prices have been matched by reductions in media price., currently \$17 for 680 megabyte recordable disk and continuing to plummet as more vendors enter the market. Another valuable capability for a recordable media which is not supported by ISO-9660 is the ability to add information to an existing volume, called multi-session in CD terminology.

Current Compact Disc-Write Once (CD-WO) technology has emerged since 1988. CD-WO is an evolution of CD-ROM to a sequential write once medium and is defined by the Orange Book; it supports writability and updatability. The Orange Book which was developed by N. V. Philips and Sony specifies the medium and the basic CD-WO system. Consequently, a new volume and file structure standard was needed as ISO 9660:88 was needed for CD-ROM. One of the differences that the CD-WO standard has that makes it different from ISO 9660:88 is that it has logical visibility to the tracks and sessions on a CD, as tracks and sessions are made visible through the use of track records in the CD-WO standard. The goals in drafting the standard were as follows:

1. support for the CD-WO write functionality,
2. overcome the deficiencies of ISO 9660:88,

- 3, support for future extensions (such as a support for the Windows NT), and
4. maintain compatibility with ISO 9660:88 in the new framework.

The fields used in the internal structure of the Descriptors are similar in format and value to the fields used in the internal format of ISO 9660:88. Changes to various fields have been made in the CD-WO structures to facilitate the support to the Portable operating System Interface (POSIX). This allows the standard to encompass the dominant operating systems in the market such as DOS, MacOS, UNIX, and VMS. Participants of Frankfurt Group were well represented by these operating system developers.

Moreover, there will be other volume and file structure standards for different media such as WORM and rewritable by the time when the CD-WO standard becomes the international standard, ISO 13490. Thus, when someone mounts a physical volume, operating systems have to know what type of standards is recorded on that medium. In order to do that, ISO 13346 (Volume and File Structure standard for non-sequential write-once media and rewritable media), DIS 13490, and possibly new tape standard have created a common volume recognition scheme. These commonly defined volume recognition sequences will enable systems to mount the correct volume standard and to boot from the medium.

Further, the defined character sets have been expanded from the ISO 9660:88, and the need for special separator characters has been reduced. The Extended Attribute Record as defined in ISO 9660:88 has been eliminated and the functionality of extended attributes has been expanded through the use of Extended Attribute area in the Directory and Path 'table' Records.

This standard has followed the same path as ISO 9660. When the High Sierra group drafted a proposal for CD-ROM volume and file structure standard, ECMA (i. e. European Computer Manufacturer's Association) modified and accepted it as ECMA 119. It was then accepted as ISO 9660 from Draft International Standard (DIS) 9660. Similarly, Frankfurt group has drafted a proposal for CD-ROM and CD-WO volume and file structure standard which ECMA modified and accepted it as ECMA 168. Right now, DIS 13490 is in international review. It is very likely to be accepted as ISO 13490 following the review period (November 93).

This standard is organized into four parts:

"Part 1: General" states notations and definitions used in this standard.

"Part 2: Volume and boot block recognition" describes volume and boot recognition structures for interchange with other standards.

"Part 3: Volume and file structure" describes volume and file descriptors along with a new, more efficient extended attribute capability.

"Part 4: Record structure" defines the various records types such as fixed record, variable record, stream record, etc.

All parts of the standard are independent: One can implement only parts 1 & 2. for volume recognition and boot purposes. 'till, Once a system recognize which driver to mount, the system can mount its native volume

and file system. If one implements parts 1, 2, and 3 for a volume and file structure without the record structure support, a transportable volume for interchange with other operating systems can be created. Part 4 can be implemented for systems which support record structures such as VMS. There will be a lot of implementations which support only parts 1, 2 and 3; for example, Mac[IS, DOS, and UNIX where record structures don't exist. The market will determine what parts will need to be supported. In the next section, a detailed description of the CD-WO volume and file structure standard is explained.

There is another standard called Rock Ridge. Its intent is to use CD-ROM as a complete implementation of X/Open and POSIX file system and directories. The purpose of the Rock Ridge initiative is to create and agree upon a common format by utilizing the system area in the directory record of the ISO 1990:88 while maintaining compatibility with the installed base of ISO 1990:88 hardware and software. The System Use Area (SUA) in the directory record includes necessary information such as UID (user id), GID (Group id), mode bits, major and minor device, UID and GID numbers by receiving systems for X/Open systems. Most UNIX CD-ROMs already use this standard specially for interchanges among UNIX systems. As the Rock Ridge standard is needed for ISO 9660:88, another standard will be needed to support future operating systems such as Windows NT with DIS13490. Thus, there are a lot of ways to implement system specific needs embedded in the descriptors, and system and application use areas along with the registration process.

It is stated that conforming DIS13490 receiving systems (i.e. systems which read data on CD-WO) shall support a subset of ISO 9660:88; the Extended Attribute Record (XAR) in ISO 9660:88 won't be supported by receiving systems. Thus, it is possible to have a disc which has both ISO 9660 & DIS13490 on a same medium. This scheme will be described in a later section. This new standard will provide an efficient architecture for faster access of directories and files compared to ISO 9660:88.

II. The Structure of DIS 13490

In this section, key Descriptors used in DIS 13490 are illustrated. The Descriptors can be categorized by three main parts -

CD-WO Extended Area:

A CD-WO Extended Area is a container for volume and file descriptors (figure 1). Volume Descriptor Sets (VDS) are located in what is called a CD-WO Extended Area. The Volume Descriptor Set contains at least one Primary Volume Descriptor (PVD) and Terminating Descriptor (TD), and zero or more Supplementary Volume Descriptor (SVD), End Transaction Descriptor (ETD), and Volume Partition Descriptor (VPD).

A Primary Volume Descriptor identifies the volume, the volume set to which the volume belongs, attributes of the volume, the character sets used in recording the contents of certain fields within the Primary Volume Descriptor, and the rule for recording and locating the End Transaction Descriptor. The End Transaction Descriptor is described in the next section.

An Supplementary Volume Descriptor provides an alternate identification of the volume and the volume set to which it belongs. A Volume Partition Descriptor specifies a partition of the Volume, attributes of the partition and an identification of the partition. The Terminating Descriptor ('1'1)) identifies the termination of a Volume Descriptor Set. Because CD-WO is an updatable medium, the volume information can be updated by writing new Volume Descriptor Sets. The standard also specifies how to recognize the most recent or valid Volume Descriptor Set.

Also, in a CD-WO Extended Area or at a location where it is pointed to by the End Transaction Descriptor, File System Descriptor Set (FSDS) is there to identify: the file set, the set of characters allowed in certain fields of the descriptors associated with the file set, attributes of the file set, options] application and publisher information, and optional copyright, abstract, and bibliographic information. A File System Descriptor set contains File Set Descriptor (FSD - a concept of a logical volume), and Implementation Use Descriptor (IUD). This File System Descriptor Set is directly associated with a path table which identifies every directory in the directory hierarchy describing the set of files in the file set. Thus at least one of File System Descriptor Set should be written over a volume to describe the content of a disc. Also, zero or more Implementation Use Descriptor (IUD) can be there to identify an implementation and contain information for that implementation's use. This is one of descriptors which allow support for current and future operating systems' needs that is not defined by this standard.

in a CD-WO Extended Area, different Boot Descriptors (BD) can be written. This is designed to allow a generic boot ROM. Such a boot ROM would scan for Boot Descriptors with a matching system identification (which might represent combinations of processor type and memory management). After examining the boot identifier, which might encode the operating system type and options, the boot ROM can present the operator with a choice of operating systems to boot.

All Volume Descriptor Sets, File System Descriptor Sets, and Boot Descriptors are enclosed in a Beginning Extended Area Descriptor (BEAD) and Terminating Extended Area Descriptor (TEAD). "1'11[1s, the CD-WO extended area identifies that CD-WO volume and file structure standard was used to write this disc. Figure 1 demonstrates a minimum required descriptors for a volume 01 a volume set. The Boot Descriptor is optional.

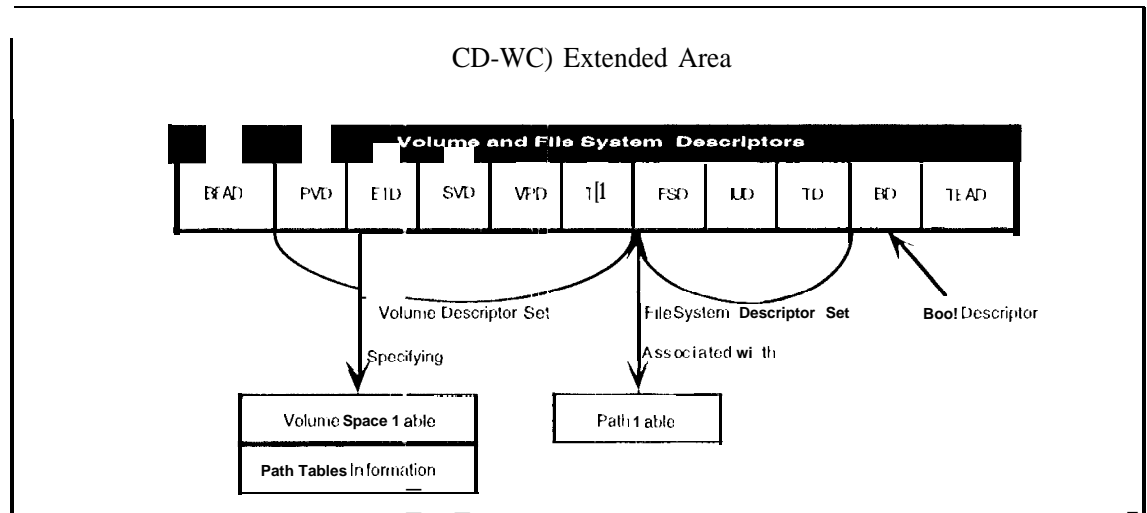


Figure 1: A Minimum Descriptors needed in DIS 13490 CD-WC Extended Area.

End Transaction Descriptor (ETD):

This is one of the most important descriptors in DIS 13490. As the information on a disc is updated, Volume Descriptor Sets, File System Descriptor Sets, and End Transaction Descriptors are used to update the volume and file structure. The more a disk is updated, the more complicated these sequences of descriptors are. It has, this standard suggests several rules for recording End Transaction Descriptors in order to search for ETDs effectively. An End Transaction Descriptor contains a pointer to the current and previous Volume Descriptor Sets and File System Descriptor Sets, and the previous ETD sequence. ETDs contain two arrays. The first is the Path Table Information (PTI) and the second is Volume Space Table which consists of the Track Specification Record (TSR). The Path Table Information contains records which points to the location of the Path Table or the directory of path tables that correspond to a particular character set that is identified in a corresponding Primary Volume Descriptor and the file sets identified by File Set Descriptors. The Track Specification Record contains a set of records which contains information about the tracks such as size and recording format.

The figure 2 illustrates the ETD and the relationship between a PVD, and previous and current ETDs.

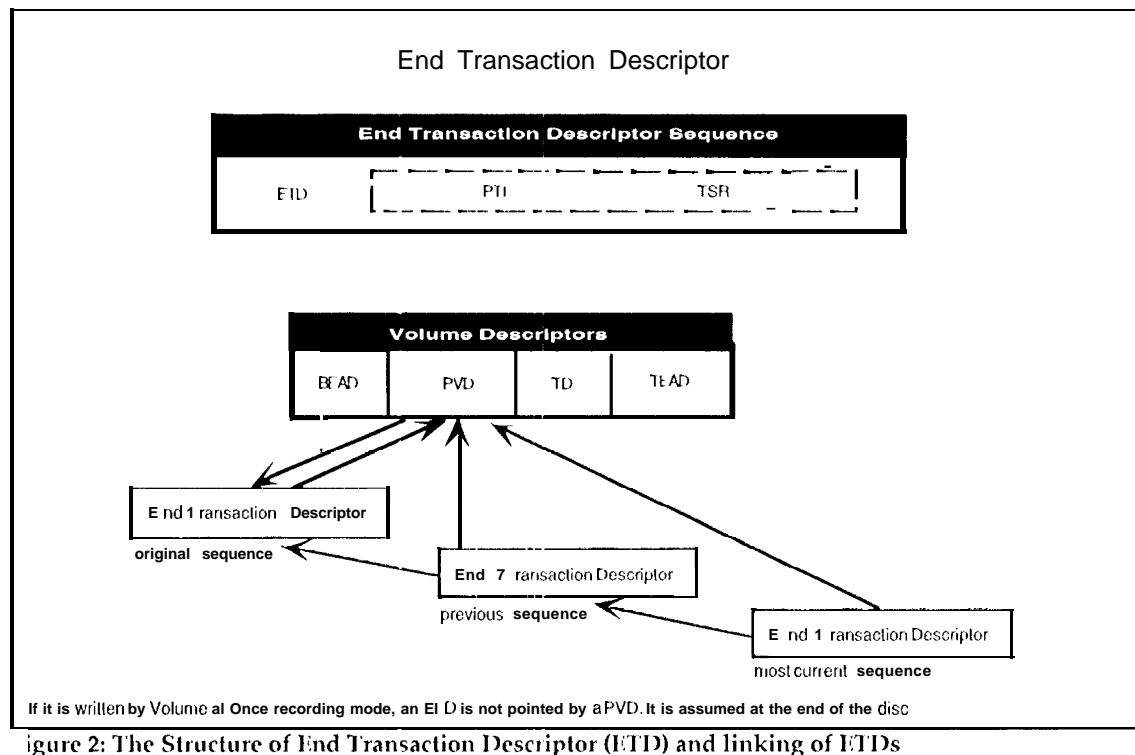


Figure 2: The Structure of End Transaction Descriptor (ETD) and linking of ETDs

Path Tables, Directories, and Extend Attribute (XA):

Multiple Primary Volume Descriptors and Path Tables information (PTI) allow multiple directory hierarchies. A Path Table specifies:

- a. the root of a directory hierarchy,
- b. each directory in the directory hierarchy, and its relationship, if any, with other directories in the directory hierarchy, and
- c. the location and size of the directory file for each directory in the directory hierarchy.

Volume structure descriptors, the Path Table Records, and the Directory Records have an Extended Attribute Area (XAA). This area is used to provide extended functionality to this standard by providing the incorporation of tag identified attributes. These attributes are associated with an ordinary file when the XAA contains tags in the directory record and are associated with a directory when the XAA contains tags in the path table record. The XAA may contain several attributes; the number of which can exceed the desired Directory (or Path Table) Record size or such a record may contain attributes deemed by the implementation to be located in another extent. In this case, the "CB" extended attribute field (XAF) is used to provide this capability as well as partitioning the continuation extent identified by the "CB" XAF into a System Use Area and an Application Use Area.

By using the descriptors mentioned above, the standard allows updating of files on CD-WO media as described in the Philips/ Sony Orange Book. Consequently, the files may be modified, the directory hierarchy may change, and additional directories may be added within the constraint that

write once media will retain all previous revisions of recorded information. The figure 3 and 4 show a simple set of directory hierarchies that can be constructed with this standard, and the updated structure when a file is added to a subdirectory.




Figure 3: Example of Directory Structure on CD-WO.

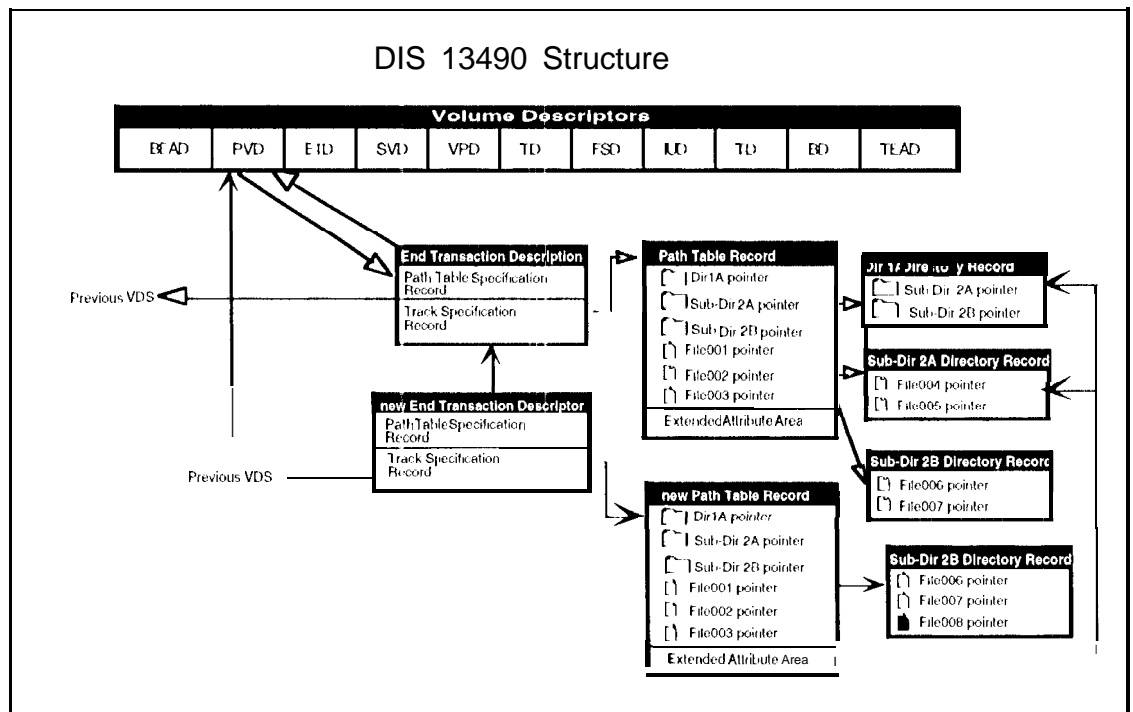


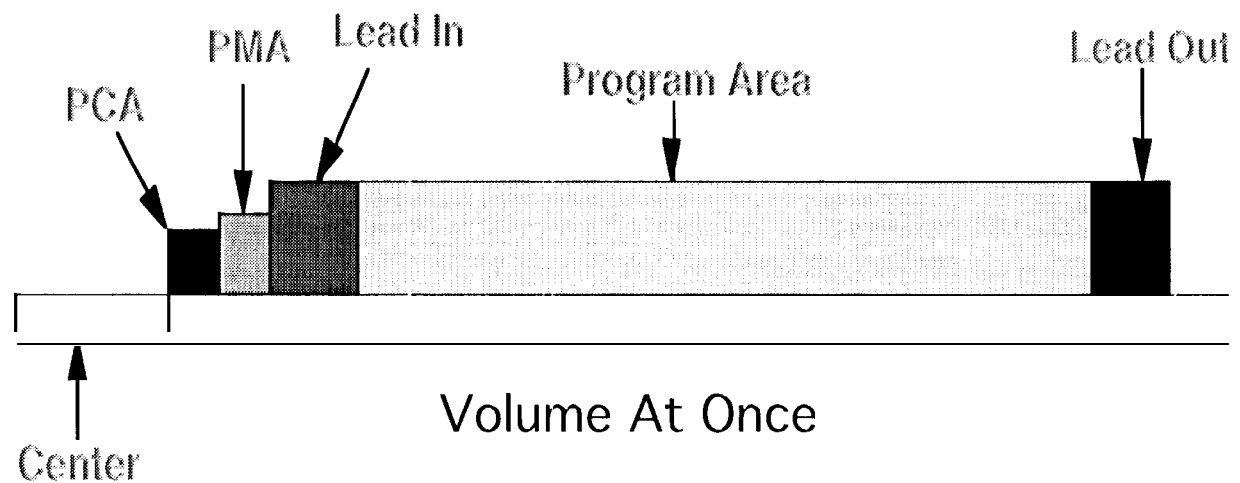


Figure 4: DIS 13490 Structure Illustrating Figure 3 Directory Structure.


111, Recording Methods

The Standard of Recording (i.e. Orange Book Part 11: CD-WO, N. V. Philips and Sony Corp., November 1990) identifies three methods which specify tracks on the disc that can be recorded. They are Track-at-Once, Incremental with Fixed length Packets, and Incremental with Variable length Packets. These methods have different addressing methods which are described below.

- 'Track-at-once' refers to recording one or multiple tracks in one uninterrupted stream. A simpler version of 'Track-at-once' is Volume-at-once which refers to recording a complete CD-WO volume in one uninterrupted stream; this is the same method how CD-ROM is mastered. By using the 'Track-at-once' method, one can do stepwise creation of a CD-WO with up to 99 tracks possible.



• Incremental recording refers to recording within a track with multiple data streams, and these streams are separated by gaps. Depending on the sizes of these streams, whether it is a fixed length or a variable length is determined. However, this scheme involves for existing players to recognize new information which reside among data streams. Thus, it can't be detected by existing CD-AUDIO and CD-ROM players.



These recording methods have their own advantages specific to certain applications. Track-at-once is the only method widely used in the current market.

A special version of the CD-WO disc is the multi-session disc. A session is a sequence of one or more tracks where the track numbers form a continuous ascending sequence. Each session could be an ISO 9660:88 volume. However, both CD-ROM drives (Yellow Book) and CD-WO/CD-R (Orange Book) drives do not support all types of recording schemes. For example, most CD-ROM drives can only read discs recorded with a single session Track-at-Once method. Also, one brand of CD-WO recorder can only record discs with a Track-at-Once method. Users need to be aware of these limitations when deciding what type of drives to purchase. Although the logical format of the disc conforms to the standard, the disc might not be readable by some drives. Within a year most of the CD-WO drives will be able to record and read all three types of recording schemes. Also, new CD-ROM drives should be able to read all types of CD-WOs.

I would like to thank Thomas Wong at SUN Microsystems and Mike Martin at JPL for their input.

Note: The research described in this article was carried out by the Jet Propulsion Laboratory, (California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, (California Institute of Technology).

Note: DIS 13490, also known as ECMA 168, is available from ECMA, European Computer Manufacturers Association, 114 Rue du Rhone - CH-1204 Geneva (Switzerland), Phone: +41 22 735 36 34, Fax: +41 22 786 52 31. For further information on CD-WO or DIS 13490, please contact Jason Hyon at Jet Propulsion Laboratory, E-Mail: jhyon@jplpds.jpl.nasa.gov or (818) 306-6054.